

IMAGE RESTORATION WITH LOCALLY SELECTED CLASS-ADAPTED MODELS

Afonso M. Teodoro, José M. Bioucas-Dias, Mário A. T. Figueiredo

Instituto de Telecomunicações and
Instituto Superior Técnico, Universidade de Lisboa,
Portugal

ABSTRACT

State-of-the-art algorithms for imaging inverse problems (namely deblurring and reconstruction) are typically iterative, involving a denoising operation as one of its steps. Using a state-of-the-art denoising method in this context is not trivial, and is the focus of current work. Recently, we have proposed to use a class-adapted denoiser (patch-based using Gaussian mixture models) in a so-called plug-and-play scheme, wherein a state-of-the-art denoiser is plugged into an iterative algorithm, leading to results that outperform the best general-purpose algorithms, when applied to an image of a known class (e.g. faces, text, brain MRI).

In this paper, we extend that approach to handle situations where the image being processed is from one of a collection of possible classes or, more importantly, contains regions of different classes. More specifically, we propose a method to locally select one of a set of class-adapted Gaussian mixture patch priors, previously estimated from clean images of those classes. Our approach may be seen as simultaneously performing segmentation and restoration, thus contributing to bridging the gap between image restoration/reconstruction and analysis.

Index Terms— Gaussian mixture models, image denoising, image deblurring, plug-and-play, class-adapted priors.

1. INTRODUCTION

Denoising is one of the oldest and most central problems in image processing, dating back to the early days of the field [13]. Although it has been argued that the current best generic methods are very close to the theoretically maximum possible performance [9], image denoising is still a very active area of research. Most, if not all, state-of-the-art methods belong to the patch-based family, i.e., they process the noisy image on a patch-by-patch fashion, relying on techniques such as non-local means [8], collaborative filtering [10], dictionary learning [2], or statistical models of the patches [17], [24], [25], [26]. In more general imaging inverse problems, such as deblurring, computed tomography, or magnetic resonance imaging, to mention only a few classical examples, it may not be obvious how these patch-based techniques may be applied, and this has recently been a topic of interest [11], [20], [22], [25].

Whereas the bulk of the work in image restoration aims at developing methods of general applicability, there has been some recent work on developing class-specific methods [15], [18]. As the name suggests, these methods are tailored to perform very well on

a certain class of images, for example, text, faces, or some type of medical images. Indeed, if we knew the class of the image being processed/reconstructed, it would be expectable that a targeted method could outperform a generic one, i.e., one that does not take into account the specific class of the image in hand. Whereas in some cases, it is known that the image being estimated belongs to a certain class (e.g., brain MR or CT images, face images, fingerprints, text), in many situations this is not the case and there may even be regions of different classes present in the image (for example, a document image may contain text and one or more faces or some other type of images). This second type of scenario (where several classes may be present in a given image) is the one addressed in this paper. Identifying the classes that are present at each location of an image can be interpreted as a form of image segmentation, thus the problem formulated and addressed in this paper may be seen as that of simultaneous image restoration and segmentation.

Recently [18], we have proposed a method for class-adapted restoration/reconstruction, which builds upon the so-called *plug-and-play* approach [19], by plugging a class-adapted denoiser based on Gaussian mixture models (GMM) into the iterations of an *alternating direction method of multipliers* (ADMM) algorithm. Experiments reported in [18] (both in deblurring and compressive imaging) have shown that the proposed method yields state-of-the-art results, when applied to images known to contain text or a face, clearly outperforming the best generic techniques, such as IDD-BM3D [11].

In this paper, we extend the plug-and-play class-adapted approach to the scenario mentioned above: the image being restored is from an unknown class and/or may contain regions from different classes (e.g., text and faces, or text and natural images). As mentioned above, this may be seen as a method to perform simultaneous segmentation and restoration, by exploiting synergies between these two tasks. Notice that our main goal is not to obtain a good or meaningful segmentation, but to use the segmentation to allow class-specific models to be exploited at each location of the image, thus our focus is on the restoration performance of the method. Nevertheless, we believe that this type of approach may contribute to bridging the gap between low-level (restoration) and mid-level (segmentation) image processing/analysis.

The remaining sections of the paper are organized as follows. Section 2 reviews the classical image restoration formulation and the basic tools upon which our method is built: ADMM, the plug-and-play scheme, and patch-based image denoising using GMM priors. The proposed method is described in Section 3, while experimental results are reported in Section 4. Finally, Section 5 concludes the paper and gives pointers for future work.

This work was partially supported by the (Portuguese) *Fundação para a Ciência e Tecnologia* (FCT), grants UID/EEA/5008/2013 and BD/102715/2014.

2. FORMULATION AND TOOLS

2.1. Problem Formulation

The classical formulation of image reconstruction/restoration problems is

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^m$, $\mathbf{x} \in \mathbb{R}^n$ are the observed data and the underlying clean vectorized image, respectively, $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the observation matrix, and \mathbf{n} is noise. For simplicity, we assume that the noise is Gaussian, with zero mean and known variance σ^2 . In denoising, the operator \mathbf{A} is the identity, but in more general cases it is non-invertible or ill-conditioned, making (1) an ill-posed inverse problem. Arguably, the most common approach to handle ill-posed problems of the form (1) is to seek a *maximum a posteriori* (MAP) estimate,

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2\sigma^2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \phi(\mathbf{x}), \quad (2)$$

where $\frac{1}{2\sigma^2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 = -\log p(\mathbf{y}|\mathbf{x}) + A$ (with A an irrelevant constant) is the negative log-likelihood and $\phi(\mathbf{x}) \propto -\log p(\mathbf{x})$ is the negative log-prior. Some recent work has also been devoted to obtaining *minimum mean squared error* (MMSE) estimates, rather than MAP, which corresponds to the posterior expectation $\mathbb{E}[\mathbf{x}|\mathbf{y}]$ [12], [14]. The criterion in (2) may also be given a non-probabilistic interpretation, and seen as a regularization approach.

2.2. ADMM and Plug-and-Play

In recent years, variable splitting algorithms, such as ADMM, have received a lot of attention, in particular in the image processing and machine learning communities [5]. A notable feature of ADMM, as applied to (2), is that it separates the handling of the data term (log-likelihood) from that of the prior/regularizer. The standard instantiation of ADMM to tackle (2) consists in the cyclic application of the following steps [1], [3]:

$$\mathbf{x}^{k+1} := \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \frac{\mu}{2} \|\mathbf{x} - \mathbf{v}^k - \mathbf{d}^k\|_2^2, \quad (3)$$

$$\mathbf{v}^{k+1} := \underset{\mathbf{v}}{\operatorname{argmin}} \phi(\mathbf{v}) + \frac{\mu}{2} \|\mathbf{x}^{k+1} - \mathbf{v} - \mathbf{d}^k\|_2^2, \quad (4)$$

$$\mathbf{d}^{k+1} := \mathbf{d}^k - (\mathbf{x}^{k+1} - \mathbf{v}^{k+1}). \quad (5)$$

Problem (3) is quadratic and has a closed-form solution, which requires solving a linear system (equivalently, inverting a matrix):

$$\mathbf{x}^{k+1} = (\mathbf{A}^T \mathbf{A} + \mu \mathbf{I})^{-1} ((\mathbf{A}^T \mathbf{y} + \mu (\mathbf{v}^k + \mathbf{d}^k))). \quad (6)$$

Although this is sometimes seen as an obstacle (motivating, for example, the introduction of linearized versions of ADMM [23]), there are cases in which this inversion can be computed very efficiently using fast transforms (namely the FFT) [1], [3], and in those cases ADMM exhibits state-of-the-art speed. Image deconvolution (with periodic [1] or other boundary conditions [3]) and inpainting are two of those cases where ADMM excels.

The update expression in (4) corresponds to the so-called *Moreau proximity operator* (MPO) of ϕ [4], computed at $\mathbf{x}^{k+1} - \mathbf{d}^k$. Recall that the MPO of $\phi : \mathbb{R} \rightarrow \bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$ is defined as

$$\operatorname{prox}_\phi(\mathbf{z}) = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 + \phi(\mathbf{x}),$$

and can be seen as the MAP solution of a denoising problem, where the argument of prox_ϕ is the noisy data, the noise is Gaussian i.i.d.

with unit variance, and the prior is $p(\mathbf{x}) \propto \exp(-\phi(\mathbf{x}))$.

In recent work, it has been suggested that the denoiser corresponding to the MPO in (4) can be replaced with an off-the-shelf state-of-the-art denoising algorithm, such as BM3D [10]; that approach has been termed *plug-and-play* [19]. Since, in general, a denoising algorithm may not correspond necessarily to the MPO of a convex regularizer/log-prior, convergence of the resulting algorithm requires further analysis; initial results have been presented in [19].

2.3. Patch-Based Denoising with GMM Priors

It has been shown that a simple GMM, estimated from a collection of clean images, is a remarkably effective patch prior [25]. More recently, we have shown that excellent denoising performance is also obtained if this GMM is estimated directly from the patches of the noisy image to be denoised, and that the corresponding expectation-maximization (EM) algorithm is a simple modification of the one for estimating the GMM from noiseless patches [17]. Moreover, since the prior is a GMM, it is easy to compute the conditional expectation of each patch given its noisy version, which is the optimal MMSE estimate (in contrast with [25], where a MAP estimate is used). Letting \mathbf{x}_i and \mathbf{y}_i denote an arbitrary patch of images \mathbf{x} and \mathbf{y} (unknown clean image and noisy one, respectively), the probabilistic model for patch denoising is

$$p(\mathbf{x}_i) = \sum_{m=1}^K \alpha_m \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_m, \mathbf{C}_m) \quad (7)$$

$$p(\mathbf{y}_i|\mathbf{x}_i) = \mathcal{N}(\mathbf{y}_i; \mathbf{x}_i, \sigma^2 \mathbf{I}) \quad (8)$$

where $\mathcal{N}(\cdot; \boldsymbol{\mu}, \mathbf{C})$ denotes a Gaussian probability density function of mean $\boldsymbol{\mu}$ and covariance \mathbf{C} . The resulting MMSE estimate of \mathbf{x}_i is given by

$$\hat{\mathbf{x}}_i = \sum_{m=1}^K \beta_m(\mathbf{y}_i) \mathbf{v}_m(\mathbf{y}_i), \quad (9)$$

where

$$\mathbf{v}_m(\mathbf{y}_i) = (\sigma^2 \mathbf{C}_m + \mathbf{I})^{-1} (\sigma^2 \mathbf{C}_m^{-1} \boldsymbol{\mu}_m + \mathbf{y}_i), \quad (10)$$

and

$$\beta_m(\mathbf{y}_i) = \frac{\alpha_m \mathcal{N}(\mathbf{y}_i; \boldsymbol{\mu}_m, \mathbf{C}_m + \sigma^2 \mathbf{I})}{\sum_{j=1}^K \alpha_j \mathcal{N}(\mathbf{y}_i; \boldsymbol{\mu}_j, \mathbf{C}_j + \sigma^2 \mathbf{I})}. \quad (11)$$

Notice that $\beta_m(\mathbf{y}_i)$ is simply the posterior probability that the i -th patch was generated by the m -th GMM component, whereas $\mathbf{v}_m(\mathbf{y}_i)$ is the conditional MMSE (and MAP) estimate of \mathbf{x}_i , if we knew that it had been generated by the m -th GMM component.

The final denoised image is assembled by putting the patch estimates back in their locations. Since the patches overlap, there are several estimates of each pixel, which are usually combined by straight averaging. In [17], we proposed to use the optimal weighted averaging, where the weights are the inverses of the posterior variances, which can also be computed in closed-form. The use of weights based on the posterior variance of the patch estimates was also used in [9], but with a single Gaussian per patch.

Recently [18], we have proposed to use GMM patch-based denoising in a plug-and-play approach, to deblur images of specific classes. For that purpose, the GMM is estimated from a collection of clean images of the class of interest, and the denoising step (4) of the ADMM algorithm is replaced with the GMM-patch-based method described above. Experiments reported in [18] show that the method produces excellent results in deconvolving text and face images, out-

performing the generic state-of-the-art method IDD-BM3D [11].

3. PROPOSED METHOD

In this paper, we extend the method proposed in [18] to handle images of unknown classes or even containing regions from different classes. Rather than a single class-adapted GMM, estimated from a collection of clean images from that class, consider C classes, each of which modelled by a GMM,

$$p(\mathbf{x}_i | c_i) = \sum_{m=1}^{K(c_i)} \alpha_m^{(c_i)} \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_m^{(c_i)}, \mathbf{C}_m^{(c_i)}), \quad (12)$$

where $c_i \in \{1, \dots, C\}$ is the class label of the i -th patch, and C the total number of classes. To estimate the i -th patch, we begin by classifying it into one of the classes, and then use the corresponding GMM to obtain an MMSE estimate of that patch (given by (9)), conditioned on its noisy version. We emphasize that this process is repeated until some stopping criterion is met.

As mentioned in Section 1, classifying each patch into one of the classes can be seen as performing image segmentation. However, we stress again that segmentation is not the main goal of the proposed approach, thus we will only focus on its performance in terms of denoising and deblurring. To classify the patches, we consider two alternatives:

- Simply classifying each patch independently using the maximum-likelihood criterion,

$$\hat{c}_i = \arg \max_{c \in \{1, \dots, C\}} p(\mathbf{x}_i | c). \quad (13)$$

- Jointly classifying all the patches under a Markov random field prior $p(\mathbf{c})$ (where \mathbf{c} denotes the field of all the patch class labels), more specifically a Potts prior [7],

$$\hat{\mathbf{c}} = \arg \max_{\mathbf{c} \in \{1, \dots, C\}^I} \log p(\mathbf{c}) + \sum_i \log p(\mathbf{x}_i | c_i). \quad (14)$$

To solve (14) we use the α -expansion graph-cut algorithm proposed in [6]. For more details about MRF priors for image segmentation, see [6], [7].

In summary, the multi-class denoiser described in the two preceding paragraphs is used in the plug-and-play approach described in Subsection 2.2 (see Algorithm 1).

Algorithm 1: Proposed ADMM-GMM with classification

Input: Blurred image, blur kernel, generic GMM, class-specific GMMs;

Output: Deblurred image, patch classification;

Parameters: Patch size, μ ;

```

1 repeat
2   Solve (3);
3   Extract deblurred image patches;
4   Classify each patch - (13) or (14);
5   Solve (4)
6   Denoise each patch using class-adapted GMM (9);
7   Combine all estimates of each image pixel;
8   Update dual variable (5);
9 until stopping criterion;
```

4. EXPERIMENTAL RESULTS

We start by presenting some results on image denoising. Table 1 compares the results obtained with the denoising algorithm, with and without classification of the patches. As a baseline, we also present the results of a state-of-the-art denoising algorithm, BM3D [10] (with default parameters). In every run, the patch size was set to 8 by 8, and we trained a GMM with 20 components on the noisy patches, using the approach described in [17]. Furthermore, the following classes were considered: *text*, *faces*, *brain MRI*, *fingerprints*, and *generic*. All of the external GMMs, also with 20 components each, were trained with samples from the corresponding class, except the *generic* GMM, which was trained using random images from the Berkeley dataset for image segmentation (BSDS300) [16].

We conclude that the proposed modification does not have a significant impact on the denoising performance. Arguably, this is due to the fact that, in pure denoising, it is possible to estimate a GMM from the noisy image itself [17]. The resulting model is thus better adapted to the input image than if it would be trained from a different set of images, even if these images are from the same class.

Figure 1 illustrates the difference in the patch labelling, if we consider classifying each patch independently via the maximum-likelihood (ML) criterion or using the α -expansion graph-cut algorithm. While we do not expect the latter labelling to perform better in terms of denoising or deblurring, it is more meaningful from a segmentation viewpoint, since it exhibits higher spatial coherence. In this example, we used three different models: one trained from the noisy image itself (black), one targeted to text images (grey), and one trained with *generic* images (white).

In our deblurring experiments, we considered the same classes, where each GMM (with 20 components) was trained using 6×6 patches. Table 2 shows the results of ADMM-GMM algorithm, with and without classification, for all blur kernels in [11]. As a benchmark, we also present the results obtained with state-of-the-art IDD-BM3D [11] (with default parameters). In this set up, we assumed that the class of the input image is unknown a priori, even in the text or face images so, when no classification is done, the algorithm uses only the *generic* GMM. Otherwise, we let the algorithm decide which class should be used for each patch, which explains the discrepancy that exists relative to the results reported in [18]. Furthermore, after 100 iterations of the algorithm, we switched the *generic* GMM with another GMM trained from the deblurred patches, which we assume to be reasonable estimates of the clean patches by then.

Since we cannot learn a GMM from the blurred input image, the improvement that is achieved with the proposed method is much more visible when, in fact, the input image contains one or more of the considered classes. On the one hand, regarding the Camera-man and House images (from the generic class), we observe that the proposed scheme performs worse than IDD-BM3D (0.76dB in the worst-case scenario), while using classification may or may not improve the performance. On the other hand, for the remaining examples, ADMM-GMM achieves better results than the generic IDD-BM3D. Although the parameters of IDD-BM3D were not tuned for these examples, we emphasize that when comparing ADMM-GMM with and without classification, the classification scheme that was proposed in this paper consistently improves the results on the images that contain one or more of the considered classes.

Figure 2 shows the patch labelling at different iterations. At first, as one would expect, the classes are not correctly identified; as the algorithm progresses, the labelling becomes more accurate. In this example we used six different classes but, for visualization purposes, we display only three: *face* (white), *text* (grey), and *other* (black).

σ	Cameraman				House				Text			
	BM3D	GMM	GMM (C)	GMM (α)	BM3D	GMM	GMM (C)	GMM (α)	BM3D	GMM	GMM (C)	GMM (α)
5	38.29	38.37	38.39	38.38	39.83	39.89	39.88	39.87	39.01	40.15	40.50	40.38
15	31.91	31.94	32.00	32.00	34.94	34.78	34.77	34.76	30.35	31.56	31.79	31.84
30	28.64	28.46	28.51	28.50	32.09	31.83	31.84	31.83	25.04	26.25	26.61	26.70
50	26.12	26.14	26.17	26.17	29.69	29.42	29.43	29.38	20.45	22.65	22.96	22.81
100	23.07	22.97	22.99	22.99	25.87	25.90	25.90	25.89	16.19	17.75	17.95	17.43

σ	Face				Cameraman + Text				Face + Text			
	BM3D	GMM	GMM (C)	GMM (α)	BM3D	GMM	GMM (C)	GMM (α)	BM3D	GMM	GMM (C)	GMM (α)
5	39.70	39.81	39.97	39.84	38.46	38.59	38.70	38.68	35.39	36.01	35.82	35.86
15	33.38	33.34	33.47	33.14	31.58	31.77	31.91	31.95	28.78	29.50	29.47	29.47
30	29.36	29.20	29.24	29.15	27.80	27.84	28.06	28.04	23.82	24.83	24.82	24.86
50	26.79	26.19	26.31	26.30	24.52	25.10	25.37	25.34	19.96	21.58	21.63	21.59
100	22.96	22.33	22.16	22.45	21.05	21.59	21.14	21.18	16.46	17.36	17.44	17.39

Table 1. PSNR on image denoising - Methods: BM3D [10]; GMM denoiser [17]; GMM denoiser with patch classification (C); GMM denoiser with α -expansion (α).

Experiment	Cameraman						House					
	1	2	3	4	5	6	1	2	3	4	5	6
BSNR	31.87	25.85	40.00	18.53	29.19	17.76	29.16	23.14	40.00	15.99	26.61	15.15
Input PSNR	22.23	22.16	20.76	24.62	23.36	29.82	25.61	25.46	24.11	28.06	27.81	29.98
IDD-BM3D	8.85	7.12	10.45	3.98	4.31	4.89	9.95	8.55	12.89	5.79	5.74	7.13
ADMM-GMM	8.39	6.36	9.80	3.47	4.16	4.88	9.66	8.22	12.43	5.50	5.42	6.82
ADMM-GMM (C)	8.54	6.44	9.82	3.49	4.20	4.79	9.55	8.06	12.24	5.26	5.25	6.75
ADMM-GMM (α)	8.49	6.36	9.76	3.51	4.19	4.82	9.62	7.90	12.28	5.10	5.23	6.83

Experiment	Text						Face					
	1	2	3	4	5	6	1	2	3	4	5	6
BSNR	26.07	20.05	40.00	15.95	24.78	18.11	28.28	22.26	40.00	15.89	26.22	15.37
Input PSNR	14.14	14.13	12.13	16.83	14.48	28.73	25.61	22.54	20.71	26.49	24.79	30.03
IDD-BM3D	11.97	8.91	16.29	5.88	6.81	4.87	13.66	11.16	14.96	7.31	10.33	6.18
ADMM-GMM	15.28	11.52	20.84	8.65	10.56	5.74	14.51	11.95	16.54	8.29	10.58	6.05
ADMM-GMM (C)	15.69	12.06	21.58	9.07	10.93	6.48	14.69	12.23	17.01	8.38	11.12	6.94
ADMM-GMM (α)	15.64	12.00	21.37	9.06	10.86	6.33	14.58	12.59	16.56	8.32	10.78	6.88

Experiment	Cameraman + Text						Face + Text					
	1	2	3	4	5	6	1	2	3	4	5	6
BSNR	32.98	26.95	40.00	19.71	30.29	19.18	23.51	22.26	40.00	17.30	27.48	18.14
Input PSNR	18.87	18.84	17.15	21.61	19.65	29.59	15.38	22.54	14.52	16.84	15.72	28.49
IDD-BM3D	10.83	8.71	12.39	5.40	6.27	4.89	9.71	6.57	11.24	2.43	1.79	4.28
ADMM-GMM	11.41	9.11	12.98	6.25	7.74	5.21	10.71	7.55	12.56	3.97	2.88	4.10
ADMM-GMM (C)	11.55	9.32	13.25	6.29	7.86	5.26	10.98	7.79	12.59	4.13	2.93	4.21
ADMM-GMM (α)	11.51	9.29	13.21	6.28	7.75	5.24	10.98	7.77	12.58	4.14	2.91	4.18

Table 2. ISNR on image deblurring - Methods: IDD-BM3D [11]; ADMM with GMM prior [18]; ADMM with GMM prior and patch classification (C); ADMM with GMM prior and α -expansion (α).

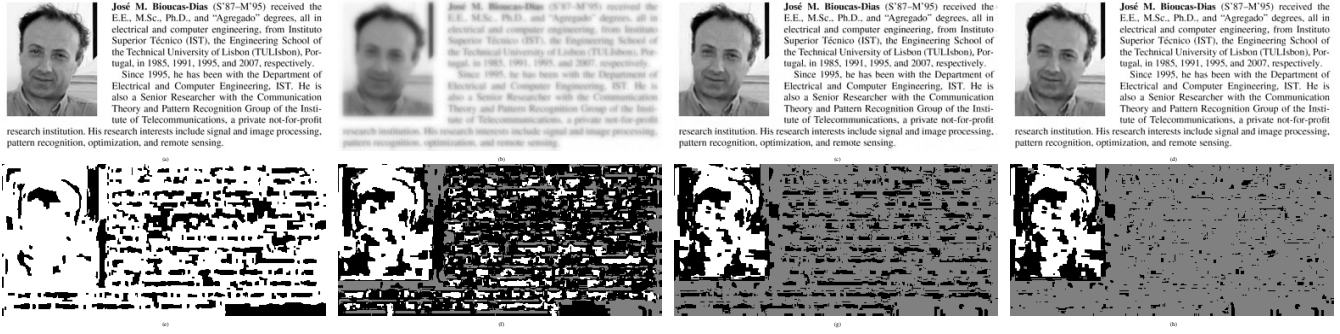


Fig. 2. Deblurring: (a) input image Face + Text; (b) blurred image (Experiment 1); (c) deblurred with IDD-BM3D [11] (ISNR = 9.71dB); (d) deblurred with ADMM-GMM with α -expansion (ISNR = 10.63dB); (e) α -expansion (1st iteration); (f) α -expansion (10th iteration); (g) α -expansion (50th iteration); (h) α -expansion (100th iteration).

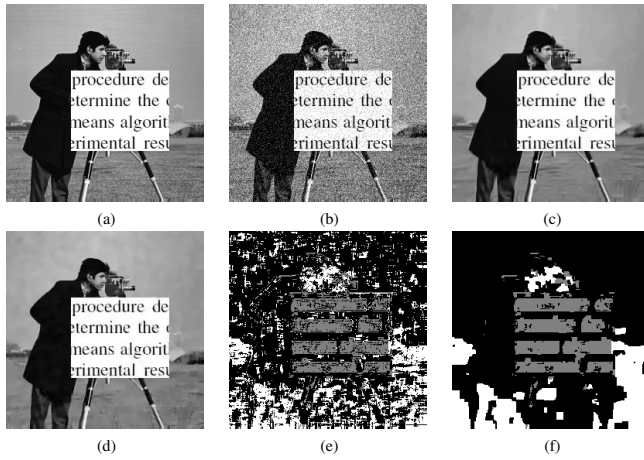


Fig. 1. Denoising: (a) input image Cameraman + Text; (b) noisy image ($\sigma = 30$); (c) denoised with BM3D [10] (PSNR = 27.80dB); (d) denoised with GMM with α -expansion (PSNR = 28.10dB); (e) ML patch classification (f) patch classification with α -expansion.

5. CONCLUSIONS AND FUTURE WORK

Recent work [15], [18] has shown that class-adapted image priors are able to outperform generic ones, when the input image in fact contains one or more of the considered classes. In this paper, we developed a method that automatically identifies which patches of the observed image should be reconstructed using class-adapted priors. This is an important feature for two main reasons: first, we often do not know whether the input image is from a particular class or not; second, we may have more than one class in a single image.

Several aspects of the proposed approach still need improvement. First, although there is some very recent work regarding the convergence of ADMM with plug-and-play denoisers, there is a need to carefully analyse the convergence of ADMM with the GMM-based denoiser. Indeed, convergence of the algorithm is observed in practice, but theoretical support is still not available. Second, in the experiments reported in Section 4, parameter μ was hand-tuned; in future work, we will pursue more sophisticated techniques, such as [21], to adjust this parameter in an automatic way.

Finally, in this paper we tested only a few very distinct image classes, but using more pre-computed models could eventually lead

to better results. Some examples of possible classes are *textures*, *sky*, *buildings*, and so on.

6. REFERENCES

- [1] M. Afonso, J. Bioucas-Dias, and M. Figueiredo, “Fast image recovery using variable splitting and constrained optimization,” *IEEE Transactions on Image Processing*, vol. 19, pp. 2345–2356, 2010.
- [2] M. Aharon, M. Elad, and A. Bruckstein, “K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation,” *IEEE Transactions on Signal Processing*, vol. 11, pp. 4311–4322, 2006.
- [3] M. Almeida and M. Figueiredo, “Deconvolving images with unknown boundaries using the alternating direction method of multipliers,” *IEEE Transactions on Image Processing*, vol. 22, pp. 3084–3096, 2013.
- [4] H. Bauschke and P. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, 2011.
- [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, pp. 1–122, 2011.
- [6] Y. Boykov, O. Veksler, and R. Zabih, “Fast approximate energy minimization via graph cuts,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, pp. 1222–1239, 2001.
- [7] Y. Boykov, O. Veksler, and R. Zabih, “Markov random fields with efficient approximations,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 648–655, 1998.
- [8] A. Buades, B. Coll, and J. Morel, “A non local algorithm for image denoising,” *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, pp. 60–65, 2005.
- [9] P. Chatterjee and P. Milanfar, “Is denoising dead?”, *IEEE Transactions on Image Processing*, vol. 19, pp. 895–911, 2010.
- [10] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising by sparse 3D transform-domain collaborative filtering,” *IEEE Transactions on Image Processing*, vol. 16, pp. 2080–2095, 2007.
- [11] A. Danielyan, V. Katkovnik, and K. Egiazarian, “BM3D frames and variational image deblurring,” *IEEE Transactions on Image Processing*, vol. 21, pp. 1715–1728, 2012.

- [12] A. Kazerouni, U. Kamilov, E. Bostan, M. Unser, “Bayesian Denoising: From MAP to MMSE Using Consistent Cycle Spinning”, *IEEE Signal Processing Letters*, vol. 20, pp. 249–252, 2013.
- [13] A. Lev, S. Zucker, and A. Rosenfeld, “Iterative enhancement of noisy images”, *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 7, pp. 435–442, 1977.
- [14] C. Louchet and L. Moisan, “Posterior expectation of the total variation model: properties and experiments”, *SIAM Journal on Imaging Sciences*, vol. 6, pp. 2640–2684, 2013.
- [15] E. Luo, S. Chan, and T. Nguyen, “Adaptive image denoising by targeted databases”, *IEEE Transactions on Image Processing*, vol. 24, pp. 2167–2181, 2015.
- [16] D. Martin, C. Fowlkes, D. Tal, and J. Malik, “A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics”, *International Conference on Computer Vision (ICCV)*, vol. 2, pp. 416–423, 2001.
- [17] A. Teodoro, M. Almeida, and M. Figueiredo, “Single-frame image denoising and inpainting using Gaussian mixtures”, *International Conference on Pattern Recognition Applications and Methods (ICPRAM)*, pp. 283–288, 2015.
- [18] A. Teodoro, J. Bioucas-Dias, and M. Figueiredo, “Image restoration and reconstruction using variable splitting and class-adapted image priors,” *IEEE International Conference on Image Processing (ICIP)*, 2016.
- [19] S. Venkatakrishnan, C. Bouman, E. Chu, and B. Wohlberg, “Plug-and-play priors for model based reconstruction”, *IEEE Global Conference on Signal and Information Processing*, pp. 945–948, 2013.
- [20] Q. Xu, H. Yu, X. Mou, D. Zhang, J. Hsieh, and G. Wang, “Low-dose X-ray CT reconstruction via dictionary learning”, *IEEE Transactions on Medical Imaging*, vol. 31, pp. 1682–1697, 2012.
- [21] Z. Xu, M. Figueiredo, and T. Goldstein, “Adaptive ADMM with spectral penalty parameter selection” *preprint arXiv:1605.07246*, 2016.
- [22] Z. Yang and M. Jacob, “Nonlocal regularization of inverse problems: a unified variational framework”, *IEEE Transactions on Image Processing*, vol. 22, pp. 3192–3203, 2013.
- [23] Z.-Z. Yang and Z. Yang, “Fast linearized alternating direction method of multipliers for the augmented ℓ_1 -regularized problem”, *Signal, Image and Video Processing*, vol. 9, pp. 1601–1612, 2015.
- [24] G. Yu, G. Sapiro, and S. Mallat, “Solving inverse problems with piecewise linear estimators: From Gaussian mixture models to structured sparsity,” *IEEE Transactions on Image Processing*, vol. 21, no. 5, pp. 2481–2499, 2012.
- [25] D. Zoran and Y. Weiss, “From learning models of natural image patches to whole image restoration”, *International Conference on Computer Vision (ICCV)* pp. 479–486, 2011.
- [26] D. Zoran and Y. Weiss, “Natural images, Gaussian mixtures and dead leaves,” *Advances in Neural Information Processing Systems (NIPS)*, vol. 25, pp. 1736–1744, 2012.